

Alleviating Overfitting for Polysemous Words for Word Representation Estimation Using Lexicons

Yuanzhi Ke

Graduate School of Science and Technology
Keio University
Yokohama, Japan
Email: enshi@soft.ics.keio.ac.jp

Masafumi Hagiwara

Graduate School of Science and Technology
Keio University
Yokohama, Japan
Email: hagiwara@soft.ics.keio.ac.jp

Abstract—Though there are some works on improving distributed word representations using lexicons, the improper overfitting of the words that have multiple meanings is a remaining issue deteriorating the learning when lexicons are used, which needs to be solved. An alternative method is to allocate a vector per sense instead a vector per word. However, the word representations estimated in the former way are not as easy to use as the latter one. Our previous work uses a probabilistic method to alleviate the overfitting, but it is not robust with small corpus. In this paper, we propose a new neural network to estimate distributed word representations using a lexicon and a corpus. We add a lexicon layer in continuous bag-of-words model, and a threshold node after the output of the lexicon layer. The threshold rejects the “bad” outputs of the lexicon layer that are less likely to be the same with their inputs. In this way, it alleviates the overfitting of the polysemous words. The proposed neural network can be trained using negative sampling, which maximizing the log probabilities of target words given the context words, by distinguishing the target words from random noises. We compare the proposed neural network with continuous bag-of-words model, the other works improving it, and the previous works estimating distributed word representations using both a lexicon and a corpus. The experimental results show that the proposed neural network is more efficient and balanced for both semantic tasks and syntactic tasks than the previous works, and robust to the size of the corpus.

I. INTRODUCTION

Natural language processing is still a challenging research area of artificial intelligence. Especially, it is a difficult issue to recognize and represent the implicit features of a piece of text properly.

Distributed text representations estimated using a neural network are useful to be applied to conventional natural language processing algorithms [1]–[11]. Great improvement by the distributed text representations estimated this way has been reported in name entity recognition and chunking [12], text classification [13]–[16], topic extraction [17], [18], and machine translation [19], [20] etc.

However, some natural language processing tasks are still challenging. For example, the conventional algorithms fail to correctly predicate the number of starts of 40% amazon reviews [16]. It indicates the needs of higher quality text representation to improve the conventional algorithms.

The early approaches to estimate text representations use n-gram models [1], [4], [5]. Mikolov et al. propose continuous bag-of-words and skip-gram models [7], [8]. Their

method outperforms the previous algorithms and costs less time. Pennington et al. [10] propose an algorithm using both local information and global information in the corpus and report a higher performance. Bojanowski et al. [11] extend the models of Mikolov et al. using the character-level information of words. Their reported experimental results outperform the original models in syntactic tasks, but fail to achieve an obvious improvement in semantic tasks.

Lexicons are useful for us humans to learn a language. We can use them to help machines to learn natural languages as well. Chen et al. [21] use the definitions in the lexicons to estimate representations for word senses and outperform the sense representations by Huang et al. [5]. Other researches take advantage of the defined synonyms or paraphrases in lexicons. Yu et al. [22] and Bollegala et al. [23] estimate the word representations by not only maximizing the probability of target word given a context, but also minimizing the distance of the paraphrases in a lexicon at the same time. Faruqui et al. [24] propose a method refining trained word representation vectors using lexicons. Xu et al. [25] estimate the word representations jointly by minimizing the distance of the tail word from the sum of the vectors of the head word and the relation for a triplet of words (*head*, *relation*, *tail*), and making words less similar to each other in a larger category. However, even though the previous methods using the paraphrases in lexicons reported improvements in syntactic analogical tasks, all of them failed to outperform the baselines in semantic analogical tasks.

There is a not well addressed issue of the previous researches on using paraphrases in lexicons to improve estimated distributed word representations: For polysemous words that have different synonyms in different contexts, if we use the paraphrases to represent them without disambiguation, they may be over-fitted to improper senses. However, disambiguation is another difficult issue. Besides, it is less easy to use if a word has several vectors for its different senses instead of one vector per word, because the usage of such word representations in the conventional systems requires additional word disambiguation. Ke et al. [26] propose a method considering the lexicon as a fuzzy set of paraphrases and using Bernoulli distribution subjected to the membership function of paraphrases to alleviate the problem. Although the method outperforms the previous works in a large corpus, the experimental results show that it is weak at small corpora.

In this paper, we propose a new method to improve the distributed word representations using paraphrases in a lexicon




Fig. 1: Continuous Bag-of-Words (CBOW) Model.

that is able to alleviate the overfitting of the polysemous words without disambiguation. Our method is efficient and easy to be combined with the conventional algorithms estimating word representations using corpora. In the experiment, our method outperforms the previous methods using paraphrases and keeps efficient for corpora in different sizes.

II. RELATED WORKS

A. Continuous Bag-of-Words

Continuous Bag-of-Words (CBOW) with negative sampling [8] is an efficient algorithm to estimate distributed word representations. The objective of CBOW is to maximize the log probability of a target word given the vectors of the context words. Denote the size of the vocabulary as V , the size of the word vectors as W . The model is like Fig. 1.

Negative Sampling is an efficient method to maximize the log probability. It is a simplified Noise Contrastive Estimation [27]. It trains the model by distinguishing target from randomly drawn noise. Denote the vector of target word as v_{w_O} , the vector of a context word as v_{w_I} , for each context word of the target word, the objective is to maximize:

$$\log \sigma(v_{w_O}^T v_{w_I}) + \sum_{i=1}^n E_{w_i \sim P_n(w)} [\log \sigma(-v_{w_i}^T v_{w_I})]. \quad (1)$$

Here, $P_n(w)$ is the distribution of the noise. σ is a sigmoid function, $\sigma(x) = 1/(1 + e^{-x})$.

B. Continuous Bag of Fuzzy Paraphrases

Continuous Bag of Fuzzy Paraphrases (CBOFP) [26] proposed by Ke et al. is a model based on CBOW to learn word representations using both a lexicon and a corpus. It is able to alleviate the overfitting of the word vectors for polysemous words. It outperforms the previous works using lexicons to estimate distributed word representations, but is not robust for small corpora. Fig. 2 shows the structure of the model.

CBOFP adds a lexicon layer to CBOW. Unlike the previous work using lexicons to estimate distributed word representations, in CBOFP, every paraphrase is a fuzzy member of the paraphrase set of a word with a degree of truth. The outputs of the lexicon layer are dropped out randomly. The dropout is controlled by a function of the paraphrases' degrees of truth that returns 0 or 1 drawn from a Bernoulli distribution. Denote

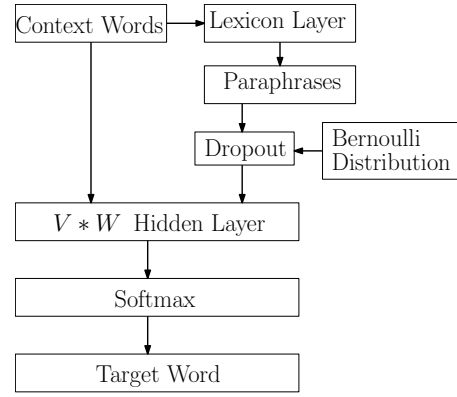


Fig. 2: Continuous Bag of Fuzzy Paraphrases (CBOFP).

the degree of truth as x , the control function $f(x)$ is defined as the following:

$$f(x) \sim \text{Bernoulli}(x). \quad (2)$$

The degree of truth is measured using the score provided by a paraphrase database called PPDB [28]–[30]. Denote the score as S , the paraphrases set as L , the degree x is calculated as the following:

$$x = \frac{S}{\max_L S}. \quad (3)$$

The reported experimental results show that this method outperforms the previous works, especially in semantic tasks, but not robust with small corpora such as text8¹.

III. THE PROPOSED METHOD

A. Structure

Our previous work CBOFP is weak at small corpus because it involves a probabilistic method to alleviate the overfitting of polysemous words. Such a method requires enough amount of training data. Thus we consider another method not using a probabilistic way.

Instead of dropping out some of the outputs of the lexicon layer randomly, we add a node after the lexicon layer as shown in Fig. 3. The inputs are the context word of the target word. They are both input into the hidden layer that contains the word vectors to learn, and input into the lexicon layer. The lexicon layer outputs the paraphrases of the inputs. The node after the lexicon layer takes the score of the paraphrase and holds a threshold. If the score of the paraphrase is higher than the threshold, it returns true. Otherwise, it returns false. The paraphrases whose outputs are false are not learned. True paraphrases of the context words are input into the hidden layer to learn together with the original context words. The output of the hidden layer is a vector of the vocabulary size. After softmax, the neural network outputs the target word. By maximizing the probability of outputting correct word, the

¹<http://mattmahoney.net/dc/text8.zip>

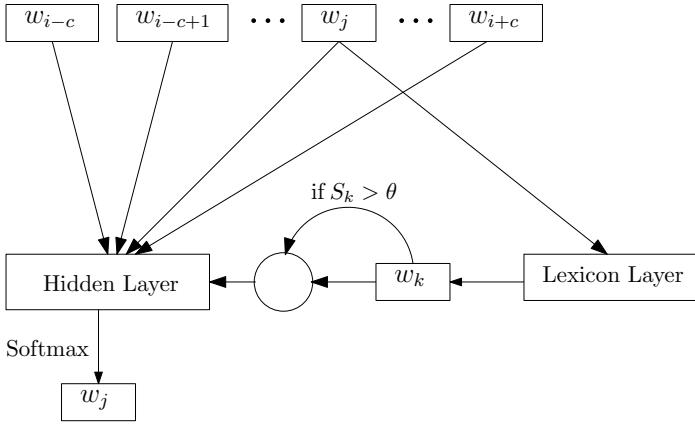


Fig. 3: The proposed neural network.

hidden layer can be trained and used as the vectors represent the words in the vocabulary.

B. The Lexicon Layer

We use a paraphrase database called PPDB2.0 [28]–[30] to build our lexicon layer, which has been used in the previous works [22], [24], [26]. The paraphrases in it are extracted automatically from multilingual resources. It is reported useful in many other tasks such as recognizing textual entailment [31], [32], measuring the semantic similarity [33]–[35], monolingual alignment [36], [37], and natural language generation [38].

PPDB2.0 provides not only the paraphrases, but also the features, the alignment types and the entailment types. There are six types of entailment in PPDB2.0 as shown in Table I. We do not consider the “exclusion” type, “other related” type, “independent” type, and only use the more “natural” paraphrases of the “equivalence” type, “forward entailment” type, and “reverse entailment” type.

For each word, we build a vector representing the set of its paraphrases in the lexicon layer. Let us denote each paraphrase in PPDB2.0 as *(headword, tailword, relationship)*. For word *A, B, C, D, E*, if there are paraphrases *(A, B, Equivalence)*, *(A, C, Forward Entailment)*, *(D, A, Reverse Entailment)*, *(A, E, Exclusion)*, *(A, F, Independent)*, and *(A, G, Other Related)*, the paraphrase set vector for *A* is *(B, C, D)*, and *E, F, G* are discarded.

As the score input to threshold node, we use the PPDB2.0 scores. They are estimated by a supervised scoring model on

TABLE I: Different types of relationships of paraphrases in PPDB2.0 [29], [30].

Relationship	Description
Equivalence	X is the same as Y
Forward Entailment	X is more specific than/is a type of Y
Reverse Entailment	X is more general than/encompasses Y
Exclusion	X is the opposite of Y / X is mutually exclusive with Y
Other Related	X is related in some other way to Y
Independent	X is not related to Y

the basis of human judgments for 26,455 paraphrase pairs and high correlation of the PPDB2.0 scores and human judgments are reported [29].

C. Learning the Word Representations

With the threshold, the objective to maximize becomes

$$\sum_{w_i \in G} \sum_{j \in C} \left[\log p(w_i | w_j) + \sum_{w_k \in L_{w_j}} f(S_{jk}) \log p(w_i | w_k) \right]. \quad (4)$$

Here, G is the set of words in the corpus, C is the context of word w_i , L_{w_j} is the paraphrase set of the context word w_j , S_{jk} is the score of paraphrase w_k in L_{w_j} . The function $f(S_{jk})$ is defined as the following:

$$f(S_{jk}) \begin{cases} 1 & \text{if } S_{jk} > \theta \\ 0 & \text{if } S_{jk} < \theta. \end{cases} \quad (5)$$

Here, θ is the threshold of the threshold node.

The log probability is maximized in the learning phrase using negative sampling. Similarly to that for CBOW, we maximize the log probability in equation (4) by maximizing

$$\log \sigma(v_{w_i}^T v_{w_j}) + \sum_{n=1}^N E_{w_i \sim P_n(w)} [\log \sigma(-v_{w_n}^T v_{w_j})], \quad (6)$$

$w_n \neq w_i, w_n \notin L_{w_j}.$

We draw noise that does not equal the target word or is not in the paraphrase set of the target word from the noise distribution $P_n(w)$. The target of the noise is labeled with zero, and the neural network is trained by maximizing the probability of the target word, given the input from the input layer and the lexicon layer, while minimizing the probability of the target word, given the noise, at the same time.

IV. EXPERIMENTS

A. The Corpus Used in the Experiments

We use text8 and a larger corpus called enwiki9² for the experiments because we are to evaluate if the proposed model is more robust for a smaller corpus than the previous work and keeps efficient for a larger one. Both text8 and enwiki9 are part of the English Wikipedia³ dump. Text8 contains 16,718,843 tokens while enwiki9 contains 123,353,508 tokens. The vocabulary size of text8 is 71,291, while that of enwiki9 is 218,317. We see that text8 is one tenth the size of enwiki9.

²<http://mattmahoney.net/dc/enwiki9.zip>

³<https://en.wikipedia.org/>

TABLE II: Examples of the questions in the word analogical reasoning task.

Semantic	Beijing	China	Tokyo	Japan
	boy	girl	brother	sister
Syntactic	amazing	amazingly	calm	calmly
	aware	unaware	clear	unclear

B. The Task for Evaluation

The word analogical reasoning task introduced by Mikolov et al.[8] is used for evaluation in the experiments. For a quater-nion of words (w_A, w_B, w_C, w_D) in which the relationship of w_A and w_B is similar to that of w_C and w_D , the objective is to predict w_D on the basis of w_A, w_B and w_C by searching the word whose vector is the closest to $v_B - v_A + v_C$. The dataset has a semantic part and a syntactic part. In the semantic part, (w_A, w_B) and (w_C, w_D) have a similar semantic relationship while they have a syntactic one in the syntactic part. Table II shows an example of the questions in the task.

There are 8,869 questions in the semantic part and 10,675 questions in the syntactic part.

C. Tuning

To find the proper value of threshold θ , we let the proposed neural network learn the word representations for text8 and enwiki9 and run the word analogical reasoning task with different threshold. The results for text8 is shown in Fig. 4. The results for enwiki9 is shown in Fig. 5.

We see that the correlation of the performance and the threshold is not linear. And it is not very same for different tasks or corpora. However, because all of the score of the paraphrases —the PPDB2.0 scores are less than seven in the version we used, we can find the best threshold in the interval.

From Fig. 4, we see that the best threshold for the semantic part and the total dataset is 3.8, using text8. The best threshold for the syntactic part is 5.7. From Fig. 5, we see that the best threshold for the semantic part and the total dataset is 1.5, using text9. The best threshold for the syntactic part is 5.2. It shows that lower threshold is better for a larger corpus and semantic tasks. The best accuracies are shown in Table III.

The other parameters are set as those used by the public word2vec demo⁴ that are already well tuned. The initial learning rate is set to 0.05. The number of negative samples drawn in negative sampling is set to 25. The context window is set to 8. The total iteration time is set to 25. And the size of the word vectors in the hidden layer is set to 200.

D. Comparison and Results

We compare our proposed neural network with CBOW [8], CBOW enriched with subword information [11], GloVe [10], the work of Faruqui et al. [24], jointReps [23], RC-Net [25] and CBOFP [26]. For the first four, we got their public online available implements by the authors. We used the implements by the authors to learn word representations using text8 and enwiki9 and report the results. For jointReps and RC-Net, we

TABLE III: The best accuracies in the word analogical reasoning task.

	Semantic [%]	Syntactic [%]	Total[%]
text8	48.80	47.09	44.73
enwiki9	73.62	59.63	66.13

failed to find an available implements that correctly run using text8 and enwiki9. Thus we use the results in their papers for jointReps and RC-Net. The reported results of RC-Net are achieved using enwiki9 as well, while those of jointReps are achieved using ukWaC⁵.

In Table IV, we compare our proposed neural network under text8 with CBOW, CBOW enriched with subword information, GloVe, the work of Faruqui et al. and CBOFP. We do not compare with jointReps and RC-Net here, because there are no reported results evaluated by the word analogical reasoning tasks of these methods with a similar corpus.

We can see that using text8, our proposed neural network achieves the best accuracy in the semantic, and outperforms the others except CBOW enriched with subword information in syntactic part. CBOW enriched with subword information is reported powerful for representing syntactic features, but ours is more balanced.

In Table V, we compare our proposed neural network with the previous works using enwiki9. The results of jointReps and RC-Net here are the reported results in their paper. The results of JointReps are not achieved using enwiki9 but using ukWaC.

We see that our proposed neural network outperforms the previous works for the semantic part and the whole dataset under enwiki9. The proposed neural network failed to outperform CBOW enriched with subword information in the syntactic part. However, our proposed neural network is better at the semantic part than the CBOW enriched with subword information and outperforms it in the overall accuracy.

All of the experimental results show that the proposed neural network is more balanced than the previous works, more powerful than the other works using lexicons to estimate or improve distributed word representations, benefiting from the threshold node alleviating the overfitting of polysemous words. Moreover, while the CBOFP failed to outperform most of the others under text8, the proposed neural network keeps outperform all the others in semantic parts and achieves the second best overall accuracy. It shows that the proposed neural network is more robust to the size of the corpus.

V. CONCLUSIONS

To alleviate the overfitting of the polysemous words that is not well addressed in the previous works, we proposed a new neural network estimating distributed word representations in this paper. Additional to the conventional continuous bag-of-words model, we added a lexicon layer, and a threshold node after the output of the lexicon layer. The threshold is manually tuned. The neural network can be trained using negative sampling. The experimental results show that the

⁴<https://code.google.com/archive/p/word2vec/>

⁵<http://wacky.sslmit.unibo.it>




Fig. 4: Threshold tuning for text8.




Fig. 5: Threshold tuning for text9.

TABLE IV: Comparison against the previous works using text8.

	Semantic [%]	Syntactic [%]	Total[%]
The proposed	48.80	47.09	44.73
CBOW [8]	46.72	41.90	43.91
Enriched CBOW[11]	15.95	73.62	49.63
GloVe [10]	41.15	21.59	29.73
Faruqui [24]	34.80	51.53	44.57
CBOFP [26]	46.35	42.13	43.88

TABLE V: Comparison against the previous works using text9.

	Semantic [%]	Syntactic [%]	Total[%]
The proposed	73.62	59.63	66.13
CBOW [8]	72.65	59.25	65.33
Enriched CBOW[11]	33.08	75.39	56.19
GloVe [10]	66.35	43.46	53.80
Faruqui [24]	53.88	61.31	57.94
JointReps [23]	61.46	69.33	65.76
RC-Net [25]	34.36	44.42	-
CBOFP [26]	73.29	59.44	65.85

proposed neural network is more powerful and balanced than the previous models using lexicons to estimate or improve distributed word representations. Besides, unlike our previous work CBOFP, the proposed neural network in this paper is robust to small corpora. Auto tuning of the threshold and the other parameters is a remaining issue. We are going to work on it in the future.

REFERENCES

- [1] Y. Bengio, R. Ducharme, P. Vincent, and C. Jauvin, "A neural probabilistic language model," *journal of machine learning research*, vol. 3, no. Feb, pp. 1137–1155, 2003.
- [2] A. Mnih and G. Hinton, "Three new graphical models for statistical language modelling," in *Proceedings of the 24th international conference on Machine learning*. ACM, 2007, pp. 641–648.
- [3] A. Mnih and G. E. Hinton, "A scalable hierarchical distributed language model," in *Advances in neural information processing systems*, 2009, pp. 1081–1088.
- [4] R. Collobert, J. Weston, L. Bottou, M. Karlen, K. Kavukcuoglu, and P. Kuksa, "Natural language processing (almost) from scratch," *Journal of Machine Learning Research*, vol. 12, no. Aug, pp. 2493–2537, 2011.
- [5] E. H. Huang, R. Socher, C. D. Manning, and A. Y. Ng, "Improving word representations via global context and multiple word prototypes," in *Proceedings of the 50th Annual Meeting of the Association for Computational Linguistics: Long Papers-Volume 1*. Association for Computational Linguistics, 2012, pp. 873–882.
- [6] T. Mikolov, "Statistical language models based on neural networks," Ph.D. dissertation, Brno University of Technology, 2012.
- [7] T. Mikolov, I. Sutskever, K. Chen, G. S. Corrado, and J. Dean, "Distributed representations of words and phrases and their compositionality," in *Advances in neural information processing systems*, 2013, pp. 3111–3119.
- [8] T. Mikolov, K. Chen, G. Corrado, and J. Dean, "Efficient estimation of word representations in vector space," in *ICLR Workshop*, 2013.
- [9] T. Mikolov, W.-t. Yih, and G. Zweig, "Linguistic regularities in continuous space word representations," in *Proceedings of NAACL HLT*, vol. 13, 2013, pp. 746–751.
- [10] J. Pennington, R. Socher, and C. D. Manning, "Glove: Global vectors for word representation," in *Empirical Methods in Natural Language Processing (EMNLP)*, 2014, pp. 1532–1543.
- [11] P. Bojanowski, E. Grave, A. Joulin, and T. Mikolov, "Enriching word vectors with subword information," *arXiv preprint arXiv:1607.04606*, 2016.
- [12] J. Turian, L. Ratinov, and Y. Bengio, "Word representations: a simple and general method for semi-supervised learning," in *Proceedings of the 48th annual meeting of the association for computational linguistics (ACL 2010)*. Association for Computational Linguistics, 2010, pp. 384–394.
- [13] R. Socher, B. Huval, C. D. Manning, and A. Y. Ng, "Semantic compositionality through recursive matrix-vector spaces," in *Proceedings of the 2012 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning*. Association for Computational Linguistics, 2012, pp. 1201–1211.
- [14] Q. V. Le and T. Mikolov, "Distributed representations of sentences and documents," in *the 31st International Conference on Machine Learning (ICML 2014)*, vol. 14, 2014, pp. 1188–1196.
- [15] Y. Kim, "Convolutional neural networks for sentence classification," in *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2014.
- [16] A. Joulin, E. Grave, P. Bojanowski, and T. Mikolov, "Bag of tricks for efficient text classification," *arXiv preprint arXiv:1607.01759*, 2016.
- [17] R. Das, M. Zaheer, and C. Dyer, "Gaussian lda for topic models with word embeddings," in *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*. Association for Computational Linguistics, 2015.
- [18] S. Li, T.-S. Chua, J. Zhu, and C. Miao, "Generative topic embedding: a continuous representation of documents," in *the 54th annual meeting of the Association for Computational Linguistics (ACL 2016)*. Association for Computational Linguistics, 2016.
- [19] W. Zaremba, I. Sutskever, and O. Vinyals, "Recurrent neural network regularization," *CoRR*, vol. abs/1409.2329, 2014.
- [20] I. Sutskever, O. Vinyals, and Q. V. Le, "Sequence to sequence learning with neural networks," in *Advances in Neural Information Processing Systems 27 (NIPS 2014)*, 2014.
- [21] X. Chen, Z. Liu, and M. Sun, "A unified model for word sense representation and disambiguation," in *EMNLP*. Citeseer, 2014, pp. 1025–1035.
- [22] M. Yu and M. Dredze, "Improving lexical embeddings with semantic knowledge," in *the 52nd Annual Meeting of the Association for Computational Linguistics (ACL2014)*. Association for Computational Linguistics, 2014, pp. 545–550.
- [23] D. Bollegala, A. Mohammed, T. Maehara, and K.-I. Kawarabayashi, "Joint word representation learning using a corpus and a semantic lexicon," in *Proceedings of the 30th AAAI Conference on Artificial Intelligence (AAAI'16)*, 2016.
- [24] M. Faruqui, J. Dodge, S. K. Jauhar, C. Dyer, E. Hovy, and N. A. Smith, "Retrofitting word vectors to semantic lexicons," in *Proceedings of NAACL*, 2015.
- [25] C. Xu, Y. Bai, J. Bian, B. Gao, G. Wang, X. Liu, and T.-Y. Liu, "Re-net: A general framework for incorporating knowledge into word representations," in *Proceedings of the 23rd ACM International Conference on Conference on Information and Knowledge Management*. ACM, 2014, pp. 1219–1228.
- [26] Y. Ke and M. Hagiwara, "Fuzzy paraphrases in learning word representations with a corpus and a lexicon," *ArXiv e-prints*, Nov. 2016.
- [27] M. U. Gutmann and A. Hyvärinen, "Noise-contrastive estimation of unnormalized statistical models, with applications to natural image statistics," *Journal of Machine Learning Research*, vol. 13, no. Feb, pp. 307–361, 2012.
- [28] J. Ganitkevitch, B. Van Durme, and C. Callison-Burch, "PPDB: The paraphrase database," in *Proceedings of NAACL-HLT*. Atlanta, Georgia: Association for Computational Linguistics, June 2013, pp. 758–764.
- [29] E. Pavlick, P. Rastogi, J. Ganitkevitch, B. V. Durme, and C. Callison-Burch, "Ppdb 2.0: Better paraphrase ranking, fine-grained entailment relations, word embeddings, and style classification," in *Association for*

Computational Linguistics. Beijing, China: Association for Computational Linguistics, July 2015.

- [30] E. Pavlick, J. Bos, M. Nissim, C. Beller, B. V. Durme, and C. Callison-Burch, "Adding semantics to data-driven paraphrasing," in *Association for Computational Linguistics*. Beijing, China: Association for Computational Linguistics, July 2015.
- [31] I. Beltagy, S. Roller, G. Boleda, K. Erk, and R. J. Mooney, "Utexas: Natural language semantics using distributional semantics and probabilistic logic," *SemEval 2014*, p. 796, 2014.
- [32] J. Bjerva, J. Bos, R. Van der Goot, and M. Nissim, "The meaning factory: Formal semantics for recognizing textual entailment and determining semantic similarity," *Proceedings of SemEval*, 2014.
- [33] L. Han, A. Kashyap, T. Finin, J. Mayfield, and J. Weese, "Umbebiquty-core: Semantic textual similarity systems," in *Proceedings of the Second Joint Conference on Lexical and Computational Semantics*, vol. 1, 2013, pp. 44–52.
- [34] Y. Ji and J. Eisenstein, "Discriminative improvements to distributional sentence similarity," in *EMNLP*, 2013, pp. 891–896.
- [35] M. A. Sultan, S. Bethard, and T. Sumner, "Dls@ cu: Sentence similarity from word alignment," in *Proceedings of the 8th International Workshop on Semantic Evaluation (SemEval 2014)*, 2014, pp. 241–246.
- [36] M. A. Sultan, S. Bethard, and T. Sumner, "Back to basics for monolingual alignment: Exploiting word similarity and contextual evidence," *Transactions of the Association for Computational Linguistics*, vol. 2, pp. 219–230, 2014.
- [37] X. Yao, B. Van Durme, C. Callison-Burch, and P. Clark, "Semi-markov phrase-based monolingual alignment," in *EMNLP*, 2013, pp. 590–600.
- [38] J. Ganitkevitch, C. Callison-Burch, C. Napoles, and B. Van Durme, "Learning sentential paraphrases from bilingual parallel corpora for text-to-text generation," in *Proceedings of the Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, 2011, pp. 1168–1179.